Bayesian calibration for forensic evidence reporting

Niko Brümmer and Albert Swart

[1]AGNITIO Research, South Africa

{nbrummer,aswart}@agnitio-corp.com

arXiv:1403.5997v2 [stat.ML] 25 Mar 2014

Abstract

We introduce a Bayesian solution for the problem in forensic speaker recognition, where there may be very little background material for estimating score calibration parameters. We work within the Bayesian paradigm of evidence reporting and develop a principled probabilistic treatment of the problem, which results in a Bayesian likelihood-ratio as the vehicle for reporting weight of evidence. We show in contrast, that reporting a likelihood-ratio distribution does not solve this problem. Our solution is experimentally exercised on a simulated forensic scenario, using NIST SRE'12 scores, which demonstrates a clear advantage for the proposed method compared to the traditional plugin calibration recipe.

Index Terms: forensic speaker recognition, Bayesian paradigm

1. Introduction

The problem that we address is how to use the output (score) of a speaker recognizer to make minimum-expected-cost Bayes decisions. Bayes decision theory provides an *idealized mathematical model* for the legal process in a court case, where it has to be established whether the prosecution hypothesis holds *beyond a reasonable doubt*. One interpretation of this model is that it quantifies *doubt* via posterior probabilities, while the *reasonable* threshold can be expressed via the relative costs of false acquittal and false conviction. Specifically, this model prescribes, convict if

$$\frac{P(\text{prosecution hypothesis}|\text{evidence})}{P(\text{defence hypothesis}|\text{evidence})} > \frac{C_{\text{false-convict}}}{C_{\text{false-acquit}}} \qquad (1)$$

The LHS represents the doubt, while the RHS is the reasonable threshold. Another interpretation is that it chooses the course of action (convict or acquit) which minimizes the expected cost of making a wrong decision.

The threshold could instead be motivated without explicit reference to cost. Jaynes [1] suggests: Setting the threshold at 10 000 "will mean, crudely, that on the average not more than one conviction in 10 000 will be in error; a judge who required juries to follow this rule would probably not make one false conviction in a working lifetime on the bench".

The purpose of this paper is not to debate the applicability of speaker recognition or Bayes decision theory to real court cases—for examples of such debate, see [2, 3]. Our goal is to design speaker recognizers to produce outputs with a clearly defined interpretation. This interpretation is supplied by the Bayes decision model. We work within this model and try to make our recognizer as good as possible according to this interpretation.

2. Theory

We shall restrict ourselves to the scientific problem of computing the LHS of (1), with the assumption that the legal process will take care of quantifying the RHS. We follow Balding's principle [4]: "Evidence is of value inasmuch as it alters the probability that the defendant is guilty".

In the Bayesian paradigm for forensic evidence reporting, it is recommended that the posterior odds, i.e. the LHS of (1), be factored as posterior odds = prior odds × likelihood-ratio [5, 6]. The prior odds is then moved to the RHS, where it is lumped together with the threshold as part of the legal process, reducing the scientific problem to just dealing with the likelihood-ratio. In this paper we shall eventually do so—our final results will be given in terms of likelihood-ratios. But we strongly recommend that when analysing how to compute such likelihood-ratios, one does not directly jump to working with likelihood-ratios—a habit that could easily lead to incorrect application of probability theory—see [7] and further discussion below. The safe route to deriving this kind of calculation is to first express the desired posterior and then (if possible) to factor out the likelihood-ratio. In general, a single likelihood-ratio representing the evidence cannot always be factored out—see [4, 7] for examples.

2.1. Hypothesis posteriors

In what follows we shall make use of the notation of *directed graphical models* (Bayesian networks) [8] to specify conditional independence relationships. Readers are urged to familiarize themselves with this notation and how to read independence relationships from such graphs.[1]

We consider a trial that has to determine whether a defendant is the perpetrator of a crime. The prosecution hypothesis, denoted H_1, states that the defendant and the perpetrator are the same, while the defence hypothesis, H_2, states that they are different. The evidence can be separated into two parts: *speech evidence*, denoted e, and *other evidence*, denoted $\boldsymbol{\pi}$. We assume[2] that the following can be computed (or is given):

$$\pi_i = P(H_i|\boldsymbol{\pi}), \text{where } \pi_1 + \pi_2 = 1 \qquad (2)$$

In this context we can refer to $\boldsymbol{\pi}$ as the hypothesis prior and we are not concerned about how (2) is computed.

What does concern us, is computation involving the speech evidence. In general, the speech evidence has two parts: speech samples known to be of the defendant and speech samples known to be of the perpetrator of the crime. In this work however, we shall eventually limit our scope to processed speech evidence in the form of automatic speaker recognizer scores.

[1]Observation at a node with convergent arrows (or observation at any of its descendants) induces dependency between variables linked via this node; when not observed, such nodes block dependency. Conversely, nodes with divergent or aligned arrows induce dependency when not observed; and block dependency when observed. A node is 'observed' if it appears to the right of the | in probability notation.

[2]This assumption is part of the idealized mathematical model, and is not meant to describe what happens in actual legal practice.

We shall analyse our speech evidence, e, with the help of a generative model of the form $P(e|h,\theta)$, where $h \in \{H_1, H_2\}$ and where θ denotes the parameters of the generative model. In graphical notation, the variables we have identified so far are related as:

$$\theta \rightarrow e \leftarrow h \leftarrow \boldsymbol{\pi} \tag{3}$$

If θ is given, the calculation of the posterior odds is well known:

$$\frac{P(H_1|e,\boldsymbol{\pi},\theta)}{P(H_2|e,\boldsymbol{\pi},\theta)} = \frac{P(H_1|\boldsymbol{\pi})}{P(H_2|\boldsymbol{\pi})} \times \frac{P(e|\theta,H_1)}{P(e|\theta,H_2)} \tag{4}$$

where the RHS is the product of prior odds and likelihood-ratio (LR). In real life however, θ is not given. If we have a suitable, large, supervised database of speech data available, θ can be estimated and plugged into (4) and we are done. We shall refer to this version of the likelihood-ratio as the *plugin LR*:

$$R_{\text{plug}}(e|\theta) = \frac{P(e|\theta,H_1)}{P(e|\theta,H_2)} \tag{5}$$

It is often however the case in forensic scenarios that the large supervised databases that are available (for example LDC data) could be deemed too different from the speech encountered in the forensic case at hand. Then we may have to resort to using a *small* forensic database for estimating θ. Naturally, questions arise: What is the minimum size of database that I can use? What is the effect of the database size on the end-result of my calculation?

Answers are provided by enlarging the scope of our probabilistic treatment to a so-called fully Bayesian treatment. Instead of regarding θ as given, we admit it is not given and instead assume a much weaker constraint, namely some suitable prior distribution, $P(\theta|\Pi)$, parametrized by the hyperparameters Π. We denote the (small) supervised database (or *background data*) by $D = (E, L)$, where E represents a collection of (processed) speech samples and L represents the associated labels. The graphical notation for this more complete specification reads:

$$L \rightarrow E \leftarrow \theta \rightarrow e \leftarrow h \leftarrow \boldsymbol{\pi}, \tag{6}$$

with $\Pi \rightarrow \theta$ omitted for brevity. This diagram makes explicit the assumption that the background data, E, and the evidence, e, were generated by the very same model parameters. Taking into account the independence assumptions encoded by (6), the posterior odds can be expressed as:

$$\frac{P(H_1|e,D,\boldsymbol{\pi},\Pi)}{P(H_2|e,D,\boldsymbol{\pi},\Pi)} = \frac{P(H_1|\boldsymbol{\pi})}{P(H_2|\boldsymbol{\pi})} \times \frac{P(e|D,\Pi,H_1)}{P(e|D,\Pi,H_2)} \tag{7}$$

This Bayesian posterior odds also factors into prior odds and a *Bayesian LR*, which we denote:

$$R_B(e|D) = \frac{P(e|D,\Pi,H_1)}{P(e|D,\Pi,H_2)} \tag{8}$$

2.2. The Bayesian likelihood-ratio

To compute the LR (8), we can express the numerator or denominator as:

$$\begin{aligned} P(e|D,\Pi,H_i) &= \int_\Theta P(e,\theta|D,\Pi,H_i)\,d\theta \\ &= \int_\Theta P(e|\theta,D,\Pi,H_i)P(\theta|D,\Pi,H_i)\,d\theta \\ &= \int_\Theta P(e|\theta,H_i)P(\theta|D,\Pi)\,d\theta \end{aligned} \tag{9}$$

where Θ represents the support of $P(\theta|\Pi)$ and where we have dropped irrelevant conditioning terms using the independence relationships encoded by (6). We shall refer to $P(e|\theta,H_i)$ as the *plugin likelihood* and to $P(\theta|D,\Pi)$ as the *parameter posterior*. The parameter posterior represents everything we know about the parameters after having processed the background data D. The required LR now becomes:

$$\begin{aligned} R_B(e|D) &= \frac{\int_\Theta P(e|\theta,H_1)P(\theta|D,\Pi)\,d\theta}{\int_\Theta P(e|\theta,H_2)P(\theta|D,\Pi)\,d\theta} \\ &= \frac{\big\langle P(e|\theta,H_1)\big\rangle_{P(\theta|D,\Pi)}}{\big\langle P(e|\theta,H_2)\big\rangle_{P(\theta|D,\Pi)}} \end{aligned} \tag{10}$$

which is a *ratio of expected likelihoods*, with expectations taken w.r.t. the parameter posterior.

If however, one immediately starts with the agenda of *let's calculate the likelihood-ratio*, without referring to a careful specification as in (6), it is easier to fall into the trap of doing *erroneous* calculations like taking expectations of the plugin LR or its logarithm:

$$\left\langle \frac{P(e|\theta,H_1)}{P(e|\theta,H_2)} \right\rangle_\theta \text{ or } \left\langle \log \frac{P(e|\theta,H_1)}{P(e|\theta,H_2)} \right\rangle_\theta$$

In fact, following our derivation in [9], we can make the difference between the expected plugin log-LR and the correct Bayesian LR explicit:

$$\begin{aligned} \log R_B(e|D) = &\left\langle \log R_{\text{plug}}(e|\theta) \right\rangle_\theta \\ &+ \left\langle \log \frac{P(\theta|D,e,H_2,\Pi)}{P(\theta|D,e,H_1,\Pi)} \right\rangle_\theta \end{aligned} \tag{11}$$

where $P(\theta|D,e,H_i,\Pi)$ is an *augmented parameter posterior* conditioned on D as well as one additional supervised example, e, with assumed label H_i. This relationship holds for expectations over any distribution for θ that avoids zero denominators.

Consider the following scenario. A forensic practitioner is tasked with summarizing the weight of the speech evidence, e, for use in court. He decides to provide this via $R_{\text{plug}}(e|D)$. He realizes however, that his D is small and that if he had happened to have had some other small database, say D', the value $R_{\text{plug}}(e|D')$ could have been substantially different. He therefore goes and simulates the selection of many similarly sized databases and for each such simulation computes $\log R_{\text{plug}}(e|D')$, keeping e constant, but varying D'. Armed with a collection of such values, he can now summarize his findings as $\log(\text{LR}) = \mu \pm \sigma$, where μ and σ are the mean and standard deviation of the collection of simulated values.

What is the court to do with μ and σ? How does this help the court to decide between H_1 and H_2? The court could perhaps decide, if σ is too large, that the evidence cannot be trusted and that the decision should be based solely on $\frac{P(H_1|\pi)}{P(H_2|\pi)}$. Or, the court could decide, if σ is small enough, to go ahead and use μ to represent the weight of evidence. But $\mu = \left\langle \log R_{\text{plug}}(e|\theta) \right\rangle_\theta$ where the distribution for this expectation is formed by the above sampling process. Our (11) shows that μ cannot give the correct value (except perhaps by accident), because the second term is ignored.

More generally, any procedure that attempts to represent the evidence via a probability distribution over the plugin LR (e.g. [10]) fails to enable the court to compute the posterior probabilities required by (1). Submitting LR distributions to court is therefore contrary to Balding's principle that "evidence is of

value inasmuch as it alters the probability that the defendant is guilty".

If we want to work within the constraints of Bayes decision theory, probability theory should be followed to compute the posteriors. In the case of the simple model in (6), the solution is given by (10).

2.3. Integrating out θ

Computing R_B via (10) requires the computation of integrals w.r.t. θ. For a restricted class of conjugate-exponential models this can be done in closed form (we show an example later), but in general, closed-form solutions are not available [8]. This means the integrals have to be approximated. This section provides some general advice to help to avoid gross inaccuracy with such approximation.

The first apparent problem is that the parameter posterior, $P(\theta|D,\Pi)$, which is required for (10), conceals a similar integral. Recalling $D=(E,L)$ and using (6), we have:

$$P(\theta|D,\Pi) = P(\theta|E,L,\Pi) = \frac{P(\theta|\Pi)P(E|L,\theta)}{P(E|L,\Pi)} \tag{12}$$

where the *normalizer* is:

$$P(E|L,\Pi) = \int_\Theta P(\theta|\Pi)P(E|L,\theta)\,d\theta \tag{13}$$

But, since this normalizer is independent of h and θ, it cancels in (10), giving:

$$\begin{aligned} R_B(e|D) &= \frac{\int_\Theta P(e|\theta,H_1)P(E|\theta,L)P(\theta|\Pi)\,d\theta}{\int_\Theta P(e|\theta,H_2)P(E|\theta,L)P(\theta|\Pi)\,d\theta} \\ &= \frac{\int_\Theta P(e,E|\theta,H_1,L)P(\theta|\Pi)\,d\theta}{\int_\Theta P(e,E|\theta,H_2,L)P(\theta|\Pi)\,d\theta} \\ &= \frac{P(e,E|H_1,L,\Pi)}{P(e,E|H_2,L,\Pi)} \end{aligned} \tag{14}$$

where the new numerator and denominator are just the normalizers of our previously introduced augmented parameter posteriors, $P(\theta|D,e,H_i,\Pi)$.

We see now that the real challenge with the calculation of R_B is not the normalizer for $P(\theta|D,\Pi)$, but rather calculation of two separate normalizers for each $P(\theta|D,e,H_i,\Pi)$.

A common practice for approximation in Bayesian calculations is to obtain an approximate parameter posterior, say $\tilde{P}(\theta|D,\Pi)$, on which further calculations are based. In our case that would give:

$$R_B(e|D) \approx \frac{\int_\Theta P(e|\theta,H_1)\tilde{P}(\theta|D,\Pi)\,d\theta}{\int_\Theta P(e|\theta,H_2)\tilde{P}(\theta|D,\Pi)\,d\theta} \tag{15}$$

Usually $\tilde{P}(\theta|D,\Pi)$ is chosen to be a good approximation of the true posterior near its peak, while the tails receive little attention. Unfortunately, one or both of $P(e|\theta,H_i)$, as a function of θ, could have its peak far from the accurate region of $\tilde{P}$ and could therefore effectively be sampling the approximate posterior in inaccurate regions. In this situation, using independent approximations for the numerator and denominator of (14) could be more accurate.

However, (14) is not without its own pitfalls. In a situation where E contains very many examples, we could have $P(E|\theta,L) \gg P(e|\theta,H_i)$ and then finite numerical precision could cause the numerator and denominator representations to become identical. Care should be taken to decide which approximation is best under the circumstances.

3. Bayesian calibration

Unfortunately, current generative models for speaker recognition, such as PLDA [11], are too complex [12] and perhaps at the same time still too inaccurate, to allow direct Bayesian calculation of the kind required here. That is, if we represent defendant and perpetrator speech samples as i-vectors [13], so that e is a pair of i-vectors, and use PLDA as the generative model, then an accurate calculation of R_B is intractable. In fact, PLDA does not even give accurate plugin likelihood-ratios. To obtain well-calibrated likelihood-ratios from a speaker recognizer, an intermediate step, known as *calibration* [14, 15] is required—for recent examples see [16, 17, 18, 9, 19, 20].

For our current purposes, calibration can be understood as follows. We process the original speech evidence via an automatic speaker recognizer which processes the defendant and perpetrator speech samples to output a real score. A high (more positive) score favours H_1, while a low (more negative) score favours H_2. This score is now considered to be the evidence, e for the trial at hand. Our supervised database, $D=(E,L)$, has $E = s_1,\ldots,s_n$, a collection of n scores generated by the same recognizer in response to n pairs of speech samples deemed similar the ones in the trial at hand. The labels, $L=\ell_1,\ldots,\ell_n \in \{H_1,H_2\}$, indicate whether each score was computed from samples satisfying H_1 or H_2. In summary, e,h refer to the score and hypothesis of the trial at hand, while the s_t,ℓ_t refer to scores and hypotheses of the background data.

3.1. Normal score model, with conjugate prior

Bayesian calibration presents challenges in addition to the above-mentioned integration problem.

The problem in the plugin method, of choosing a family of score distributions [19], is shared by the Bayesian method. For this paper we default to Gaussian score modelling, which will suffice to demonstrate the advantages of Bayesian calibration. Future work will examine more sophisticated score models.

The challenge of assigning the prior, $P(\theta|\Pi)$, should not be underestimated. According to Jaynes [1], priors represent the unfinished half of probability theory. For some thoughts on selecting priors for forensic problems, see [21]. Here we aim for simplicity and select a conjugate prior, which gives closed-form integrals.

We let our model parameters be $\theta=(\mu_1,\mu_2,\lambda_1,\lambda_2)$, for a Gaussian model of the form:

$$P(e|H_i,\theta) = \mathcal{N}(e|\mu_i,\lambda_i^{-1}) \tag{16}$$

where the μ_i are means and the λ_i, are precisions (inverse variances). We consider the background data to be generated iid from θ, so that:

$$P(E|L,\theta) = \prod_{i=1}^{2}\prod_{t:\ell_t=H_i}\mathcal{N}(s_t|\mu_i,\lambda_i^{-1}) \tag{17}$$

The conjugate prior is Gaussian-gammma, of the form [8]:

$$P(\theta|\mu_0,\beta,a,b) = \prod_{i=1}^{2}\mathcal{N}(\mu_i|\mu_0,\beta^{-1}\lambda^{-1})\mathcal{G}(\lambda_i|a,b) \tag{18}$$

where $\mathcal{G}$ is a gamma distribution with parameters $a,b>0$. By choosing $a=b \ll 1$, λ_i has an expected value of 1 and a very large variance, making this prior non-informative about the scale of the scores. Likewise, to be non-informative about the

score location, we can arbitrarily choose $\mu_0 = 0$ and $\beta \ll 1$, which gives very large variance for μ_i.

Finally, $R_B(e|D)$ is computed by solving (9) in closed form, which results in a T distribution—for details, see [22, 23, 8].

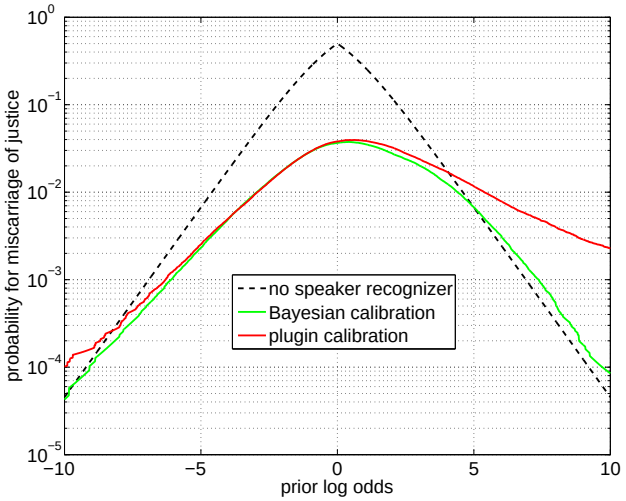

Figure 1: Bayesian vs Plugin calibration, trained on 9 H_1 and 27 H_2 examples.

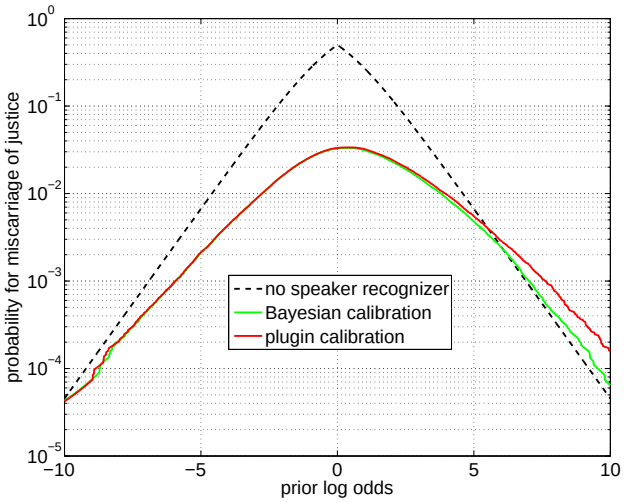

Figure 2: Bayesian vs Plugin calibration, trained on 30 H_1 and 405 H_2 examples.

4. Experiments

Our experimental setup is similar to that in [9, 19], using scores from a single speaker recognizer (an i-vector PLDA system), which was part of the ABC submission [24] to the NIST SRE'12 speaker recognition evaluation [25]. We selected background scores, D, from the ABC development database, containing pre-SRE'12 speech, while the scores for the evidence in the trial at hand, e, came from SRE'12.

We constructed D using either 3 or 10 speakers, with 3 recordings per speaker. For 3 speakers, this gave 9 H_1 scores and 27 H_2 scores, while for the 10 speaker case this gave 30 H_1 and 405 H_2 scores. The error-rates were averaged in each case over 1000 different random selections of speakers for D.

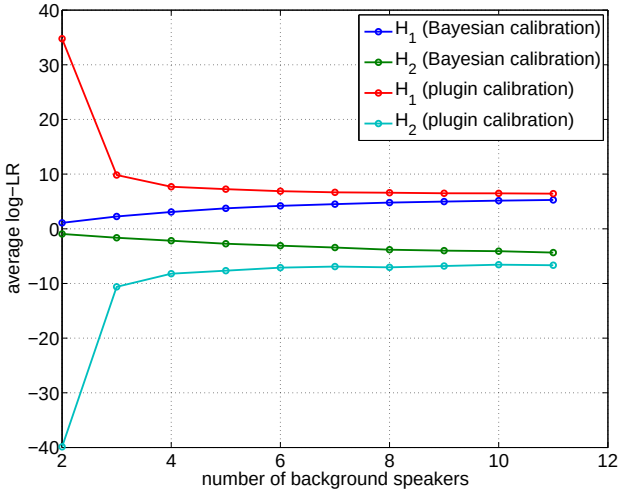

Figure 3: Plugin calibration becomes overconfident with insufficient data. Bayesian calibration moderates its confidence as a function of the data size.

To test how well it works, we vary the hypothesis prior π_1 over a wide range of values and plot the error-rate that results when using our calibrated likelihood-ratios to make decisions using (1). In a real court case the relative costs matter, but for our purposes of evaluating accuracy, we can conveniently set the costs to 1 without loss of generality [15, 26]. This error-rate can be interpreted as the probability for a miscarriage of justice (false acquit, or false convict).

The results are shown for D having 3 speakers (figure 1) and for 10 speakers (figure 2), where we compare: (i) traditional plugin calibration; (ii) the proposed Bayesian calibration; and (iii) simply convicting if $\pi_1 > \pi_2$, giving an error-rate of $\min(\pi_1, \pi_2)$. The horizontal axis is $\log \frac{\pi_1}{1-\pi_1} = \log \frac{\pi_1}{\pi_2}$.

Figure 3 compares the average hypothesis-conditional log-LR values for plugin vs Bayesian, as a function of the amount of background data. The Bayesian method behaves intuitively—and the plugin method counterintuitively.

4.1. Discussion

Both plugin and Bayesian calibrations face two main challenges on this data. The Bayesian method solves the first challenge:

(i) Limited training data. When the data is very limited (3 speakers) the Bayesian method clearly does better. With more data (10 speakers) the problem is less severe and the gap closes.

(ii) This data exhibits mild *dataset shift* [27], where the score distributions change between background and test sets. Neither method has a mechanism for dealing with dataset shift. Future work should examine ways of including dataset shift modelling into the Bayesian method.

5. Conclusion

We have shown that by widening the scope of our probabilistic treatment, the applicability of score calibration can be enlarged to include challenging situations with very limited background data. With limited data, the traditional plugin method misbehaves by being overconfident, while the Bayesian method moderates its confidence as a function of the amount of data.

6. References

[1] E. T. Jaynes, *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.

[2] J. Campbell, W. Shen, W. Campbell, R. Schwartz, J. Bonastre, and D. Matrouf, "Forensic speaker recognition," *IEEE Signal Processing Magazine*, vol. 26, 2009.

[3] P. Rose, "Where the science ends and the law begins: Numerical likelihood ratio-based forensic voice comparison in a $150 million telephone fraud," *The International Journal of Speech, Language and the Law*, vol. 20, no. 2, 2013.

[4] D. Balding, *Weight-of-Evidence for Forensic DNA Profiles*, ser. Statistics in Practice. Wiley, 2005.

[5] J. Gonzalez-Rodriguez, P. Rose, D. Ramos-Castro, D. Torre-Toledano, and J. Ortega-Garcia, "Emulating DNA: Rigorous quantification of evidential weight in transparent and testable forensic speaker recognition," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 15, no. 7, pp. 2104–2115, Sep. 2007.

[6] D. Ramos-Castro, "Forensic evaluation of the evidence using automatic speaker recognition systems," Ph.D. dissertation, Universidad Autónoma de Madrid, Madrid, Spain, Nov. 2007.

[7] N. Brümmer and E. de Villiers, "What is the 'relevant population' in Bayesian forensic inference?" AGNITIO Research, South Africa, Tech. Rep., 2011. [Online]. Available: http://arxiv.org/abs/1403.6008

[8] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer, 2007.

[9] N. Brümmer and D. Garcia-Romero, "Generative modelling for unsupervised score calibration," in *ICASSP*, 2014.

[10] G. W. Beecham and B. S. Weir, "Confidence interval of the likelihood ratio associated with mixed stain DNA evidence," *Journal of Forensic Sciences*, vol. 56, no. S1, 2011.

[11] P. Kenny, "Bayesian speaker verication with heavy-tailed priors," in *Odysey 2010: The speaker and language recognition workshop*, Brno, Czech Republic, 2010, pp. 249–252.

[12] J. Villalba and N. Brümmer, "Towards fully Bayesian speaker recognition: Integrating out the between-speaker covariance," in *Interspeech*, Florence, 2011.

[13] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech and Language Processing*, 2010.

[14] N. Brümmer and J. A. du Preez, "Application-independent evaluation of speaker detection," *Computer Speech and Language*, vol. 20, no. 2–3, pp. 230–275, 2006.

[15] N. Brümmer, "Measuring, refining and calibrating speaker and language information extracted from speech," Ph.D. dissertation, Stellenbosch University, 2010.

[16] M. Mandasari, R. Saeidi, and D. van Leeuwen, "Calibration based on duration quality measure function in noise robust speaker recognition for NIST SRE'12," in *Biometric Technologies in Forensic Science*, Nijmegen, 2013.

[17] D. van Leeuwen and N. Brümmer, "The distribution of calibrated likelihood ratios," in *Interspeech*, 2013.

[18] N. Brümmer and G. Doddington, "Likelihood-ratio calibration using prior-weighted proper scoring rules," in *Interspeech*, 2013.

[19] N. Brümmer, A. Swart, and D. van Leeuwen, "A comparison of linear and non-linear calibrations for speaker recognition," in *Odyssey 2014*, submitted.

[20] M. McLaren, A. Lawson, L. Ferrer, N. Scheffer, and Y. Lei, "Trial-based calibration for speaker recognition in unseen conditions," in *Odyssey 2014: The Speaker and Language Recognition Workshop*, Joensuu, Finland, submitted.

[21] N. Brümmer, "Tutorial for Bayesian forensic likelihood ratio," AGNITIO Research, South Africa, Tech. Rep., 2013. [Online]. Available: http://arxiv.org/abs/1304.3589

[22] ——, "Fully Bayesian score calibration assuming Gaussian distributions," AGNITIO Research, South Africa, Tech. Rep., 2011. [Online]. Available: http://niko.brummer.googlepages.com

[23] T. P. Minka, "Inferring a Gaussian distribution," MIT Media Lab, Technical Report, 1998.

[24] AGNITIO, BUT, and CRIM, "ABC SRE'12 presentation," in *NIST SRE 2012 Workshop, Orlando*, 2012.

[25] The National Institute of Standards and Technology, "The NIST year 2012 speaker recognition evaluation plan," 2012. [Online]. Available: www.nist.gov

[26] N. Brümmer and E. de Villiers, "The BOSARIS Toolkit: Theory, algorithms and code for surviving the new DCF," in *NIST SRE'11 Analysis Workshop, Atlanta*, 2011.

[27] J. Quionero-Candela, M. Sugiyama, A. Schwaighofer, and N. D. Lawrence, *Dataset Shift in Machine Learning*. MIT Press, 2009.